arXiv:1405.1332v5 [cs.DS] 1 Jun 2015

A unified framework for linear dimensionality reduction in L1

Felix Krahmer and Rachel Ward

June 3, 2015

Abstract

For a family of interpolation norms $\|\cdot\|_{1,2,s}$ on $\mathbb{R}^n$, we provide a distribution over random matrices $\Phi_s \in \mathbb{R}^{m\times n}$ parametrized by sparsity level s such that for a fixed set X of K points in $\mathbb{R}^n$, if $m \geq Cs\log(K)$ then with high probability, $\frac{1}{2}\|\boldsymbol{x}\|_{1,2,s} \leq \|\Phi_s(\boldsymbol{x})\|_1 \leq 2\|\boldsymbol{x}\|_{1,2,s}$ for all $\boldsymbol{x} \in X$. Several existing results in the literature roughly reduce to special cases of this result at different values of s: For $s = n$, $\|\boldsymbol{x}\|_{1,2,n} \equiv \|\boldsymbol{x}\|_1$ and we recover that dimension reducing linear maps can preserve the ℓ_1-norm up to a distortion proportional to the dimension reduction factor, which is known to be the best possible such result. For $s = 1$, $\|\boldsymbol{x}\|_{1,2,1} \equiv \|\boldsymbol{x}\|_2$, and we recover an ℓ_2/ℓ_1 variant of the Johnson-Lindenstrauss Lemma for Gaussian random matrices. Finally, if $\boldsymbol{x}$ is s- sparse, then $\|\boldsymbol{x}\|_{1,2,s} = \|\boldsymbol{x}\|_1$ and we recover that s-sparse vectors in ℓ_1^n embed into $\ell_1^{\mathcal{O}(s\log(n))}$ via sparse random matrix constructions.

1 Introduction

The theory for linear dimensionality reduction in Euclidean space has been the subject of much research in recent years. The celebrated *Johnson-Lindenstrauss (JL) Lemma* says that a small set of points in high-dimensional Euclidean space can be linearly embedded into a space of much lower dimension in such a way that Euclidean distances between the points are nearly preserved [20]. More specifically, given a finite set $X \subset \mathbb{R}^n$ of size $|X| = K$, there exists a linear map $\Phi : \mathbb{R}^n \to \mathbb{R}^m$ with $m = 9\varepsilon^{-2}\log(K)$ such that $(1-\varepsilon)\|\boldsymbol{x}-\boldsymbol{y}\|_2 \leq \|\Phi(\boldsymbol{x}-\boldsymbol{y})\|_2 \leq (1+\varepsilon)\|\boldsymbol{x}-\boldsymbol{y}\|_2$ for all $\boldsymbol{x}, \boldsymbol{y} \in X$. In the language of geometric embeddings, this says that K-point subsets of ℓ_2^n can be linearly embedded into ℓ_2^m with $m = \mathcal{O}(\varepsilon^{-2}\log(K))$ and distortion $1+\mathcal{O}(\varepsilon)$. Remarkably, for a fixed finite set $X \subset \mathbb{R}^n$, taking $\Phi : \mathbb{R}^n \to \mathbb{R}^m$ as a random matrix whose entries are independent and identically-distributed mean-zero Gaussian random variables will achieve such an embedding with high probability [11]. Because such a probabilistic embedding is easy to construct and is *oblivious* to the content of X, random projections have become an efficient pre-processing step for a wide range of algorithms in numerical linear algebra, compressive sensing, manifold learning, and theoretical computer science [18, 32, 15, 25, 2, 3, 35, 8, 22].

It is natural to ask about embedding results for more general ℓ_p norms, $1 \leq p \leq \infty$, e.g., a result of the form $(1-\varepsilon)\|\boldsymbol{x}\|_{\ell_p^n} \leq \|\Phi\boldsymbol{x}\|_{\ell_p^m} \leq (1+\varepsilon)\|\boldsymbol{x}\|_{\ell_p^n}$ for $m \ll n$. Unfortunately, the strong results realized by Gaussian random matrices is specific to the case $p = 2$. In fact, an embedding result of the Johnson-Lindenstrauss type is impossible for $p \neq 2$ using *any* linear embedding: as shown by Charikar and Sahai [9] for ℓ_1 and generalized by Lee, Mendel, and Naor in [23] to ℓ_p for $1 \leq p \leq \infty$, there are arbitrarily large K-point subsets X of ℓ_p such that any linear mapping $T : X \to \ell_p^m$ incurs distortion at least $D = \Omega\left((\frac{K}{m})^{|1/p-1/2|}\right)$. In particular, for ℓ_1, dimensionality reduction is possible in general only if we allow for large distortion; that is, for an arbitrary finite set X of $|X| = K$ points in $\mathbb{R}^n$, if one wishes for a map $T : \ell_1^n \to \ell_1^m$ such that $\frac{1}{\sqrt{D}}\|\boldsymbol{x}-\boldsymbol{y}\|_1 \leq \|\Phi(\boldsymbol{x}-\boldsymbol{y})\|_1 \leq \sqrt{D}\|\boldsymbol{x}-\boldsymbol{y}\|_1$ holds for all $\boldsymbol{x}, \boldsymbol{y} \in X$, then necessarily $m \geq CD^{-2}K$. This bound is tight; see [30, 33]. The ℓ_1 norm is of particular interest for several

reasons, one of which being that in high dimensions, the ℓ_1 norm is more meaningful than the ℓ_2 norm (and much more meaningful than the ℓ_p norm for p large) for inferring neighborliness in large data sets [6, 16].

Still, the lower bounds for linear dimension reduction in ℓ_1 represent a *worst-case* bound over arbitrary sets of points. Restricting attention to structured subsets of points in ℓ_1^n, much stronger statements can be made. Of particular interest is the subset of *sparse* vectors, where we recall that $\boldsymbol{x} \in \mathbb{R}^n$ is s-sparse if it has non-zero coordinates in at most s dimensions. In [9], it was shown that an arbitrary set of K s-sparse vectors can be linearly embedded into ℓ_1^m with distortion $1+\varepsilon$ once $m \geq C\varepsilon^{-2}s^2 \log K$. A uniform result over s-sparse vectors was subsequently shown by Berinde, Gilbert, Indyk, Karloff, and Strauss in [5], which we state as a proposition.

Proposition 1 (From [5]). *Fix $n, s \in \mathbb{N}$ and $\varepsilon \in (0,1)$, and fix $m \in \mathbb{N}$ satisfying $m \geq C\varepsilon^{-2}s \log n$. There exist matrices $\Phi \in \mathbb{R}^{m \times n}$ such that*

$$(1-\varepsilon)\|\boldsymbol{x}\|_{\ell_1^n} \leq \|\Phi \boldsymbol{x}\|_{\ell_1^m} \leq (1+\varepsilon)\|\boldsymbol{x}\|_{\ell_1^n} \qquad \forall \boldsymbol{x} \in \mathbb{R}^n : \ \ \|\boldsymbol{x}\|_0 \leq s. \tag{1}$$

Such a matrix Φ is said to have the 1-restricted isometry property *of order s and level ε, or 1-RIP for short.*

Explicit constructions of such matrices are *binary* and *sparse*; specifically, a sparse binary random matrix having $d = C\varepsilon^{-1}\log(n)$ ones per column and having $m \geq C'\varepsilon^{-1}sd = C''\varepsilon^{-2}s\log(n)$ rows will have the 1-RIP with high probability [5, 12]. The 1-RIP property is essentially equivalent to the combinatorial notion of *expansion* of the sparse bipartite graph underlying the measurement matrix.

1.1 Contribution of this work

The aim of this work is to initiate a unified framework for linear dimension reduction in ℓ_1. We provide a general theorem which roughly interpolates between several existing distinct results.

More specifically, we consider a family of rearrangement-invariant block ℓ_1/ℓ_2 norms $\|\cdot\|_{1,2,s}$ on $\mathbb{R}^n$ parametrized by block size $s \in [n]$, as follows. For a given $\boldsymbol{x} \in \mathbb{R}^n$, consider a partition of its support into disjoint subsets $S_1, S_2, \dots$ so that S_1 indexes the largest s elements of $\boldsymbol{x}$ in magnitude, S_2 indexes the next s largest elements, and so on, and $S_{\lceil n/s \rceil}$ may contain between zero and $s-1$ elements. Here on out, we will refer to the (not necessarily unique) decomposition $\boldsymbol{x} = (\boldsymbol{x}_{S_1}, \boldsymbol{x}_{S_2}, \dots)$ as the *s-block decreasing rearrangement* of $\boldsymbol{x}$. The norm of interest is

$$\|\boldsymbol{x}\|_{1,2,s} := \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{S_\ell}\|_1^2} \tag{2}$$

for completeness, we verify in the appendix that this indeed defines a norm. We call this an interpolation norm as on the one extreme, $s = 1$ and $\|\cdot\|_{1,2,s} = \|\cdot\|_2$; at the other extreme, $s = n$ and $\|\cdot\|_{1,2,s} = \|\cdot\|_1$.

Together with this interpolation norm, we consider random matrices $\Phi_s \in \mathbb{R}^{m \times n}$ of the form

$$\Phi_s = A_s \circ G \tag{3}$$

where $\circ$ denotes the Hadamard (entrywise) product, and

1. $G = (g_{j,k}) \in \mathbb{R}^{m \times n}$ is a random matrix populated with independent and identically distributed standard Gaussian entries, and
2. $A_s = (a_{j,k}) \in \mathbb{R}^{m \times n}$ is a random matrix populated with zeros and ones, having exactly $d = \frac{m}{s}$ ones per column, the locations of which are chosen uniformly from $[m]$ without replacement.

Our main result is as follows.

Theorem 2. *Fix $m, n \in \mathbb{N}$ and parameter $\varepsilon > 0$. Fix $s \in [n]$ such that $m/s \in \mathbb{N}$, and consider the random matrix $\Psi_s = \sqrt{\frac{2}{\pi}}\frac{s}{m}\Phi_s$. For any fixed $\boldsymbol{x} \in \mathbb{R}^n$, it holds with probability exceeding $1 - 4n\exp(-\varepsilon^2 m/8) - 2\exp\left(-\varepsilon^2\beta_0^2(m/(8s))\left(\|\boldsymbol{x}\|_{1,2,s}/\|\boldsymbol{x}\|_2\right)^2\right)$ that*

$$(.63 - \varepsilon)\|\boldsymbol{x}\|_{1,2,s} \leq \|\Psi_s \boldsymbol{x}\|_1 \leq (1.63 + \varepsilon)\|\boldsymbol{x}\|_{1,2,s}. \tag{4}$$

Remark 3. The block norm $\|\cdot\|_{1,2,s}$ has appeared previously in [13, 26], and is closely related to the so-called $\mathcal{K}$-*interpolation norm* $K(\boldsymbol{x}, t) = \sum_{j=1}^{t^2} x_j^* + t(\sum_{j>t^2}(x_j^*)^2)^{1/2}$ (where $\boldsymbol{x}^*$ denotes the decreasing rearrangement of $\boldsymbol{x}$), which is well-known in the theory of interpolation of Banach spaces [4, 17], and was used in the related context of upper and lower bounds for Rademacher sums in [28]. As shown in Proposition 9, these two norms are equivalent up to a factor of 1.63:

$$\|\boldsymbol{x}\|_{1,2,s} \leq \mathcal{K}(\mathbf{x}, \sqrt{s}) \leq 1.63\|\boldsymbol{x}\|_{1,2,s}.$$

Theorem 2 roughly interpolates between three existing results in the literature.

1. For block-size $s = 1$, the block norm $\|\cdot\|_{1,2,s}$ coincides with the Euclidean norm $\|\cdot\|_2$, and Ψ_s reduces to a properly-normalized i.i.d. Gaussian random matrix. Up to distortion $(1.63 + \varepsilon)/(.63 - \varepsilon) \approx 2.6 + \mathcal{O}(\varepsilon)$ instead of distortion $1 + \mathcal{O}(\varepsilon)$, Theorem 2 recovers a well-established ℓ_2/ℓ_1 Johnson-Lindenstrauss concentration result for Gaussian random matrices, see for instance Lemma 5.3 of [31].
2. If $\boldsymbol{x} \in \mathbb{R}^n$ is s-sparse, then $\|\boldsymbol{x}\|_{1,2,s} = \|\boldsymbol{x}\|_1$. In this case, Theorem 2 with parameter s recovers the 1-RIP embedding result from Proposition 1 for sparse vectors, albeit only for the particular s-sparse vector $\boldsymbol{x}$ and not all s-sparse vectors. Insufficient concentration of the Gaussian matrix G prevents us from passing to a uniform sparse embedding result using $\Phi_s = A_s \circ G$ at number of measurements $m = \mathcal{O}(s\log(n))$; however, as shown in Corollary 13, we do recover the result of Proposition 1 (up to distortion $(1.63 + \varepsilon)/(.63 - \varepsilon) \approx 2.6 + \mathcal{O}(\varepsilon)$ instead of distortion $1 + \mathcal{O}(\varepsilon)$) if we use for embedding the binary matrix A_s alone, rather than the composite matrix $\Phi_s = A_s \circ G$.
3. For block-size $s = n/D^2$ in Theorem 2 with fixed constant distortion $D \geq 1$, the bock norm $\|\cdot\|_{1,2,s}$ is equivalent to the $\|\cdot\|_1$ norm up to a multiplicative factor of D. In this case, Theorem 2 produces an explicit embedding which realizes, up to a factor of $\log(n)$, the best-possible dependence of $m = \Omega(n/D^2)$ dimensions necessary for embedding an arbitrary n-point set in ℓ_1 into ℓ_1^m with constant distortion D.

While we expect that the lower and upper distortion bounds $.63 - \varepsilon$ and $1.63 + \varepsilon$ can be somewhat reduced, it is not possible to bring them down to $1 - \varepsilon$ and $1 + \varepsilon$, respectively. In Proposition 15, we provide two classes of examples and show that between them there is a constant distortion factor even in the asymptotic limit.

Remark 4. It is noteworthy that, in contrast to the hardness results for linear dimension reduction in ℓ_1, it *is* possible to approximate the ℓ_1 norm of a vector $\boldsymbol{x} \in \mathbb{R}^n$ up to distortion $1 + \varepsilon$ from a linear projection $\Phi\mathbf{x}$ into $m = \mathcal{O}(\varepsilon^{-2})$ dimensions if one considers other functions of $\Phi\mathbf{x}$ besides the ℓ_1 norm. The first result of this kind, provided in [19], shows that taking Φ to have i.i.d. Cauchy-distributed entries (or i.i.d. from a p-stable distribution more generally for approximating the ℓ_p norm, $0 < p < 2$), $\|\boldsymbol{x}\|_1 = (1 \pm \varepsilon)\text{median}_k|(\Phi x)_k|$ with high probability and with $m = \mathcal{O}(\varepsilon^{-2})$. Several subsequent works [24, 21, 29, 34] provide estimators other than the median and random embedding matrices Φ other than Cauchy random matrices which are more optimized for practical implementations.

1.2 Outline of the proof of Theorem 2

The proof of Theorem 2 naturally splits in two parts. First, we treat A_s as fixed, so that G is the only source of randomness in $\Phi = A_s \circ G$. In subsection 2.1 we use concentration estimates for sums of half-normal random variables to show that $\|\Phi \boldsymbol{x}\|_1$ has subgaussian concentration about its mean with variance $\sigma^2 = \frac{s}{m}\|\boldsymbol{x}\|_2^2$. We consider G as a Gaussian random matrix because $\mathbb{E}_G\|\Phi\boldsymbol{x}\|_1$ can be calculated explicitly:

$$\mathbb{E}_G(\|\Phi\boldsymbol{x}\|_1) = \sqrt{\frac{2}{\pi}} \sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2}. \tag{5}$$

In subsection 2.2 we show, using only that $A_s = (a_{j,k}) \in \{0,1\}^{m\times n}$ has m/s ones per column,

$$\sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2} \leq \frac{1.63m}{s} \|\boldsymbol{x}\|_{1,2,s}. \tag{6}$$

In the second part of the proof, in subsection 2.3, we treat A_s as random, and show that for fixed $\boldsymbol{x} \in \mathbb{R}^n$, a lower bound is obtained with high probability with respect to the realization of A_s:

$$\sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2} \geq (.63 - \varepsilon) \frac{m}{s} \|\boldsymbol{x}\|_{1,2,s}.$$

This lower bound follows rather directly from an estimate we show for each s-sparse component $\boldsymbol{x}_{S_\ell}$ of the block decreasing rearrangement of $\boldsymbol{x}$:

$$\sum_{j=1}^{m} \left(\sum_{k\in S_\ell} a_{j,k} x_k^2 \right)^{1/2} \geq (.63 - \varepsilon) \frac{m}{s} \|\boldsymbol{x}_{S_\ell}\|_1. \tag{7}$$

In words, (7) amounts to showing that by drawing $A \in \{0,1\}^{m\times n}$ having m/s per column at random, we do not deviate much from the ideal situation where A, restricted to the columns indexed by S_ℓ, contains *exactly* one non-zero entry in each row, whence $\sum_{j=1}^{m} \left(\sum_{k\in S_\ell} a_{j,k} x_k^2 \right)^{1/2} = \frac{m}{s}\|\boldsymbol{x}_{S_\ell}\|_1$. To verify that (7) is satisfied with sufficiently high probability, we use a "balls into bins" analysis and borrow techniques from [5] and [12] used to show that a similar matrix construction satisfies the 1-RIP in Proposition 1. We note that our binary matrix construction differs from those constructions in that we use exactly $m = sd$ measurements (d is the number of ones per column), allowing us to apply the bound (6) with minimal constant 1.63.

As shown in Section 3, these ingredients can be combined to prove Theorem 2.

Remark 5. Without much additional effort, the Gaussian random matrix G in Theorem 2 can be replaced by a Bernoulli random matrix B, that is, a matrix whose entries are independent random variables which are ± 1 with equal probability. Indeed, embeddings of the form $\Phi = A_s \circ B$ are more appealing from a practical point of view as their entries are contained in $\{0, 1, -1\}$, and are thus more easily implemented and stored. Using a Bernoulli random matrix, the exact formula (5) for the conditional expectation will no longer hold, but by the Khintchine inequality (using the optimal constants provided in [14]), we still have

$$\frac{1}{\sqrt{2}} \sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2} \leq \mathbb{E}_B(\|(A_s \circ B)\boldsymbol{x}\|_1) \leq \sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2}.$$

As a result, the analog of Theorem 2 using a Bernoulli matrix will have slightly worse constants, but otherwise remains unchanged.

1.3 Notation

Throughout the paper, $C, c, C_1, \ldots$ denote absolute constants whose values may change from line to line. For integer n, we write $[n] = \{1, 2, \ldots, n\}$. Vectors are written in bold italics, e.g. $\boldsymbol{x}$, and their coordinates written in plain text, e.g. the i-th component of $\boldsymbol{x}$ is x_i. For a subset $S \subset [n]$, $\boldsymbol{x}_S$ is the vector $\boldsymbol{x}$ restricted to the elements indexed by S, and may be treated as a dense vector in $\mathbb{R}^s$ or a sparse vector in $\mathbb{R}^n$ depending on the context. Similarly, we denote by A_S the matrix A restricted to the columns indexed by S.

We denote by ℓ_p^n the space $\mathbb{R}^n$ with the ℓ_p norm $\|\boldsymbol{x}\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}, 1 \le p < \infty$, $\|\boldsymbol{x}\|_\infty = \max_i |x_i|$. When the ambient dimension is not important, we simply write ℓ_p. The number of non-zero coordinates of a vector $\boldsymbol{x}$ is denoted by $\|\boldsymbol{x}\|_0 = |\mathrm{supp}(\boldsymbol{x})|$.

We write $f(u) = \mathcal{O}(g(u))$ if and only if there exists a positive real number M and a real number u_0 such that $|f(u)| \le M|g(u)|$ for all $u > u_0$. We write $f(u) = \Omega(g(u))$ if and only if $g(u) = \mathcal{O}(f(u))$. For a function f of random variables $\boldsymbol{X} = (X_1, X_2)$, we write $\mathbb{E}_{X_1} f(\boldsymbol{X}) := \mathbb{E}\left[f(\boldsymbol{X})|X_2\right]$ for the conditional expectation, which is itself a function of the random X_2.

Finally, recall that an embedding $f : X \to Y$ of a metric space (X, d) into a metric space (Y, d') is said to have distortion $D \ge 1$ if there are constants $A, B \ge 1$ satisfying $AB \le D$ such that for all $\boldsymbol{x}, \boldsymbol{y} \in X$,

$$A^{-1} d(\boldsymbol{x}, \boldsymbol{y}) \le d'(f(\boldsymbol{x}), f(\boldsymbol{y})) \le B d(\boldsymbol{x}, \boldsymbol{y}).$$

2 Proof ingredients

The remainder of the paper is devoted to proving Theorem 2. We provide proof ingredients in this section, and we put the ingredients together in Section 3.

2.1 Concentration lemmas

In this section we treat the binary matrix $A = A_s$ as fixed, and study the concentration of $\|\Phi\boldsymbol{x}\|_1 = \|(A \circ G)\boldsymbol{x}\|_1$ around its conditional expectation $\mathbb{E}_G\|(A \circ G)\boldsymbol{x}\|_1$. The first lemma is straightforward.

Lemma 6. *Fix $\boldsymbol{x} \in \mathbb{R}^n$ and $A = (a_{j,k}) \in \{0,1\}^{m \times n}$. Let $G = (g_{j,k}) \in \mathbb{R}^{m \times n}$ have i.i.d. Gaussian entries, and consider the random matrix $A \circ G$. The expectation of $\|(A \circ G)\boldsymbol{x}\|_1$ is given by*

$$\mathbb{E}_G\|(A \circ G)\boldsymbol{x}\|_1 = \sqrt{\frac{2}{\pi}} \sum_{j=1}^m \left(\sum_{k=1}^n a_{j,k} x_k^2\right)^{1/2}.$$

Proof. The jth coordinate of $(A \circ G)\boldsymbol{x} \in \mathbb{R}^m$ can be written as $Y_j = \sum_{k=1}^n a_{j,k} g_{j,k} x_k$; this is a mean-zero Gaussian random variable with variance $\sigma_j^2 = \sum_{k=1}^n a_{j,k} x_k^2$. It follows that the random variable $|Y_j|$ is a half-normal random variable; it has mean $\mathbb{E}(|Y_j|) = \sqrt{\frac{2}{\pi}}\sigma_j = \sqrt{\frac{2}{\pi}}\left(\sum_{k=1}^n a_{j,k} x_k^2\right)^{1/2}$. The lemma follows by noting that $\mathbb{E}_G\|(A \circ G)\boldsymbol{x}\|_1 = \sum_{j=1}^m \mathbb{E}(|Y_j|)$. ☐

We now show that if $A \circ G$ is a random matrix as above, and if A has exactly d ones per column, then $\|(A \circ G)\boldsymbol{x}\|_1$ exhibits subgaussian concentrates around its expectation $\mathbb{E}_G\|(A \circ G)\boldsymbol{x}\|_1$. To do this, we will need the following lemma, whose proof uses standard arguments and can be found in the appendix.

Lemma 7. *Fix $d, m \in \mathbb{N}$ and $\alpha > 0$. Suppose that Y_i, $i = 1, 2, \ldots, m$, are independent mean-zero Gaussian random variables with variances σ_i^2 satisfying $\sum_{i=1}^m \sigma_i^2 = d\alpha^2$. Then the random*

variable $Z = \frac{1}{d}\sum_i (|Y_j| - \mathbb{E}|Y_j|)$ *satisfies*

$$\mathrm{Prob}\left[|Z| \geq \lambda\right] \leq 2\exp\left(-\frac{\lambda^2 d}{2\alpha^2}\right), \qquad \forall \lambda \geq 0.$$

Lemma 7 implies the following result about the concentration of $\|(A \circ G)\boldsymbol{x}\|_1$.

Proposition 8. *Fix* $\boldsymbol{x} \in \mathbb{R}^n$*, and fix* $A = (a_{j,k}) \in \{0,1\}^{m\times n}$ *populated with zeros and ones and having* d *ones per column. Suppose that* $G \in \mathbb{R}^{m\times n}$ *consists of i.i.d. standard Gaussian entries, and consider the random matrix* $\Phi = A \circ G$*. Then*

$$\mathbb{P}\left[\frac{|\|\Phi\boldsymbol{x}\|_1 - \mathbb{E}_G\|\Phi\boldsymbol{x}\|_1|}{d} \geq \lambda\right] \leq 2\exp\left(-\frac{\lambda^2 d}{2\|\boldsymbol{x}\|_2^2}\right), \qquad \forall \lambda \geq 0.$$

Proof. Recall via Lemma 6 that $\|\Phi x\|_1 = \sum_i |Y_i|$ where $Y_i \sim \mathcal{N}(0, \sigma_i^2)$ are independent Gaussian random variables with variances $\sigma_j^2 = \sum_{k=1}^n a_{j,k}x_k^2$. Moreover, $\sum_j \sigma_j^2 = d\|\boldsymbol{x}\|_2^2$ thanks to $A \in \{0,1\}^{m\times N}$ having exactly d ones per column. Applying Lemma 7 with $\alpha = \|\boldsymbol{x}\|_2$ gives the stated result. ☐

2.2 Analytic ingredients

Recall that $G \in \mathbb{R}^{m\times n}$ is a Gaussian random matrix. We first show that for any binary matrix $A_s \in \{0,1\}^{m\times n}$ having $d = m/s$ ones per column, it holds that $\mathbb{E}_G\|(A_s \circ G)\boldsymbol{x}\|_1 \leq 1.63d\sqrt{\frac{2}{\pi}}\|\boldsymbol{x}\|_{1,2,s}$.

Proposition 9. *Fix* $\boldsymbol{x} \in \mathbb{R}^n$ *with* s*-block decreasing rearrangement* $(\boldsymbol{x}_{S_1}, \boldsymbol{x}_{S_2}, \dots)$*. Fix* $A = (a_{j,k}) \in \{0,1\}^{m\times n}$ *having* $d = m/s$ *ones per column. Then*

$$\sum_{j=1}^m \left(\sum_{k=1}^n a_{j,k}x_k^2\right)^{1/2} \leq 1.63d\|\boldsymbol{x}\|_{1,2,s}.$$

For the proof, we will use the following *norm inequality lemma*, which was introduced in [7].

Lemma 10 (From [7])**.** *For any* $\boldsymbol{x} \in \mathbb{R}^k$*,*

$$\|\boldsymbol{x}\|_2 \leq \frac{\|\boldsymbol{x}\|_1}{\sqrt{k}} + \frac{\sqrt{k}}{4}\left(\max_{1\leq i\leq k}|x_i| - \min_{1\leq i\leq k}|x_i|\right).$$

Proof of Proposition 9. We have

$$\begin{aligned}
\sum_{j=1}^m \left(\sum_{k=1}^n a_{j,k}x_k^2\right)^{1/2} &= \sum_{j=1}^m \left(\sum_{k\in S_1} a_{j,k}x_k^2 + \sum_{\ell=2}^{R=\lceil n/s\rceil}\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} \\
&\leq \sum_{j=1}^m \left(\sum_{k\in S_1} a_{j,k}x_k^2\right)^{1/2} + \sum_{j=1}^m \left(\sum_{\ell=2}^R \sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} \\
&\leq \sum_{j=1}^m \sum_{k\in S_1} a_{j,k}|x_k| + \sum_{j=1}^m \left(\sum_{\ell=2}^R \sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} \\
&= d\|\boldsymbol{x}_{S_1}\|_1 + \sum_{j=1}^m \left(\sum_{\ell=2}^R \sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} \qquad (8)
\end{aligned}$$

Applying Hölder's inequality to the second term and using again that A has $d = m/s$ ones per column,

$$\sum_{j=1}^{m}\left(\sum_{\ell=2}^{R}\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} \le \sqrt{m}\left(\sum_{j=1}^{m}\sum_{\ell=2}^{R}\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2} = d\sqrt{s}\left(\sum_{\ell=2}^{R}\sum_{k\in S_\ell} x_k^2\right)^{1/2}. \quad (9)$$

Applying the norm inequality (Lemma 10), we have, for each block $\ell \ge 2$,

$$\begin{aligned}\sum_{k\in S_\ell} x_k^2 &\le \frac{\|\boldsymbol{x}_{S_\ell}\|_1^2}{s} + \frac{s}{16}\|\boldsymbol{x}_{S_\ell}\|_\infty^2 + \frac{1}{2}\|\boldsymbol{x}_{S_\ell}\|_1(\|\boldsymbol{x}_{S_\ell}\|_\infty - \|\boldsymbol{x}_{S_{\ell+1}}\|_\infty)\\ &\le \frac{\|\boldsymbol{x}_{S_\ell}\|_1^2}{s} + \frac{\|\boldsymbol{x}_{S_{\ell-1}}\|_1^2}{16s} + \frac{1}{2}\|\boldsymbol{x}_{S_2}\|_1(\|\boldsymbol{x}_{S_\ell}\|_\infty - \|\boldsymbol{x}_{S_{\ell+1}}\|_\infty).\end{aligned}$$

Returning to the string of inequalities (8), we continue

$$\begin{aligned}\sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k}x_k^2\right)^{1/2} &\le d\|\boldsymbol{x}_{S_1}\|_1 + \sqrt{md}\left(\sum_{\ell=2}^{R}\sum_{k\in S_\ell} x_k^2\right)^{1/2}\\ &\le d\|\boldsymbol{x}_{S_1}\|_1 + \sqrt{md}\sqrt{\sum_{\ell=2}^{R}\left(\frac{\|\boldsymbol{x}_{S_\ell}\|_1^2}{s} + \frac{\|\boldsymbol{x}_{S_{\ell-1}}\|_1^2}{16s} + \|\boldsymbol{x}_{S_2}\|_1\left(\frac{\|\boldsymbol{x}_{S_\ell}\|_\infty}{2} - \frac{\|\boldsymbol{x}_{S_{\ell+1}}\|_\infty}{2}\right)\right)}\\ &\le d\|\boldsymbol{x}_{S_1}\|_1 + d\sqrt{\sum_{\ell=2}^{R}\left(\|\boldsymbol{x}_{S_\ell}\|_1^2 + \frac{\|\boldsymbol{x}_{S_{\ell-1}}\|_1^2}{16}\right) + \frac{s\|\boldsymbol{x}_{S_2}\|_1\|\boldsymbol{x}_{S_2}\|_\infty}{2}}\\ &\le d\|\boldsymbol{x}_{S_1}\|_1 + d\sqrt{\frac{17}{16}\sum_{\ell=2}^{R}\|\boldsymbol{x}_{S_\ell}\|_1^2 + \frac{1}{16}\|\boldsymbol{x}_{S_1}\|_1^2 + \frac{1}{2}\|\boldsymbol{x}_{S_1}\|_1\|\boldsymbol{x}_{S_2}\|_1}\\ &\le d\|\boldsymbol{x}_{S_1}\|_1 + d\sqrt{\frac{17}{16}\sum_{\ell=2}^{R}\|\boldsymbol{x}_{S_\ell}\|_1^2 + \frac{5}{16}\|\boldsymbol{x}_{S_1}\|_1^2 + \frac{1}{4}\|\boldsymbol{x}_{S_2}\|_1^2}\\ &\le d\|\boldsymbol{x}_{S_1}\|_1 + d\sqrt{\frac{21}{16}\sum_{\ell=2}^{R}\|\boldsymbol{x}_{S_\ell}\|_1^2 + \frac{5}{16}\|\boldsymbol{x}_{S_1}\|_1^2}\\ &\le \sqrt{2.625}d\sqrt{\sum_{\ell=1}^{R}\|\boldsymbol{x}_{S_\ell}\|_1^2}, \quad (10)\end{aligned}$$

where the last inequality follows by applying $|a| + |b| \le \sqrt{2}\sqrt{a^2 + b^2}$ and the proposition follows by taking the bound $\sqrt{2.625} \le 1.63$. □

Remark 11. Note that (8) and (9) imply that

$$\begin{aligned}\frac{1}{d}\sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k}x_k^2\right)^{1/2} &\le \|\mathbf{x}_{S_1}\|_1 + \sqrt{s}\|\mathbf{x}_{S_1^c}\|_2\\ &:= \mathcal{K}(\mathbf{x}, \sqrt{s}) \quad (11)\end{aligned}$$

where $\mathcal{K}(\mathbf{x}, t)$ is the interpolation norm defined in Remark 3. Given that $\|\boldsymbol{x}\|_{1,2,s} \le \mathcal{K}(\mathbf{x}, \sqrt{s})$, a byproduct of the proof of Proposition 9 is that the two interpolation norms are equivalent up to a factor of 1.63:

$$\|\boldsymbol{x}\|_{1,2,s} \le \mathcal{K}(\mathbf{x}, \sqrt{s}) \le 1.63\|\boldsymbol{x}\|_{1,2,s} \quad (12)$$

2.3 Combinatorial ingredients

Consider $A_s \in \{0,1\}^{m\times n}$ a random binary matrix having $d = m/s$ ones per column. In this section we show that with high probability with respect to the realization of such a matrix, $\|A_s\boldsymbol{x}\|_1 \approx \|\boldsymbol{x}\|_1$ for all s-sparse $\boldsymbol{x} \in \mathbb{R}^n$. We begin by showing such concentration holds for a fixed s-sparse $\boldsymbol{x} \in \mathbb{R}^n$.

Proposition 12. *Fix $n, s, d \in \mathbb{N}$ and set $m = ds$. Fix a subset $S \subset [n]$ of cardinality $|S| = s$ and an ordering $\pi_1, \pi_2, \ldots, \pi_s$ of the indices in S. Draw an $m \times n$ binary random matrix $A = (a_{j,k})$ with $d = m/s$ ones per column as follows: for each column $k \in [n]$, draw d elements $\{j_1, j_2, \ldots, j_d\}$ from $[m]$ uniformly without replacement, and set $a_{j_\ell,k} = 1$. For each $\varepsilon > 0$ it holds with probability exceeding $1 - 2s\exp(-\varepsilon^2 m/2)$ that, for any s-sparse $\boldsymbol{z} \in \mathbb{R}^n$ supported on S and satisfying $|z_{\pi_1}| \geq |z_{\pi_2}| \geq \cdots \geq |z_{\pi_s}|$,*

$$d(1 - 2e^{-1} - 2\varepsilon)\|\boldsymbol{z}\|_1 \leq \|A\boldsymbol{z}\|_1 \leq d\|\boldsymbol{z}\|_1 \tag{13}$$

and

$$d(1 - e^{-1} - \varepsilon)\|\boldsymbol{z}\|_1 \leq \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k} z_k^2\right)^{1/2}.$$

Proof. As the columns of A are independent random vectors, we can assume without loss that $S = \{1, 2, \ldots, s\}$, $(\pi_1, \pi_2, \ldots, \pi_s) = (1, 2, \ldots, s)$, and $A = A_S \in \{0,1\}^{m\times s}$.

For the moment, consider the modified probability distribution over random matrices A where, for each column $k \in [s]$ we draw d elements $\Lambda_k = \{j_1, j_2, \ldots, j_d\}$ from $[m]$ uniformly *with* replacement, and set $a_{j_\ell,k} = 1$ if $j_\ell \in \Lambda_k$ (and $a_{j,\ell} = 0$ otherwise). Note that by sampling with replacement, there may be repetitions within Λ_k and so the number of ones in any particular column may be smaller than d. Still, the total number of draws with replacement is $m = ds$, d draws per each of s columns. An equivalent way to describe this process is as throwing balls into bins: a total of m balls are thrown i.i.d. into m bins, $d = m/s$ per round for s rounds; if at least one ball is thrown into bin j during round k, then $a_{j,k} = 1$; otherwise, $a_{j,k} = 0$.

Let q_k be the fraction of the m bins which remain empty after the first k rounds of this process, that is, after the first dk balls have been tossed. As the probability that any particular bin is empty at this point is equal to $(1 - \frac{1}{m})^{dk}$,

$$\mathbb{E}(q_k) = \left(1 - \frac{1}{m}\right)^{dk} \leq \exp\left(-\frac{dk}{m}\right) = \exp(-k/s). \tag{14}$$

Using the Azuma-Hoeffding inequality, it can be shown [[27], 12.19, p. 313] that the random variable q_k concentrates around its mean according to

$$\mathbb{P}\left(|q_k - \mathbb{E}(q_k)| \geq \varepsilon\sqrt{1 - [\mathbb{E}(q_k)]^2}\right) \leq 2\exp(-2\varepsilon^2 m), \qquad \forall \varepsilon > 0.$$

Since $\mathbb{E}(q_k) \leq \exp(-k/s)$, this implies in particular that, for any $\varepsilon > 0$,

$$\begin{aligned}\mathbb{P}\left(q_k \geq \exp(-k/s) + (\varepsilon/\sqrt{2})\sqrt{k/s}\right) &\leq \mathbb{P}\left(q_k \geq \exp(-k/s) + (\varepsilon/\sqrt{2})\sqrt{1 - \exp(-2k/s)}\right)\\ &\leq \mathbb{P}\left(q_k \geq \mathbb{E}(q_k) + (\varepsilon/\sqrt{2})\sqrt{1 - [\mathbb{E}(q_k)]^2}\right)\\ &\leq 2\exp(-\varepsilon^2 m) \qquad .\end{aligned}$$

Taking a union bound over $k \in [s]$,

$$\mathbb{P}\left(\forall k \in [s] : q_k \leq \exp(-k/s) + \varepsilon\sqrt{k/s}\right) \geq 1 - 2s\exp(-\varepsilon^2 m/2), \qquad \forall \varepsilon > 0. \tag{15}$$

Back to the setting where we draw d elements $\{j_1, j_2, \ldots, j_d\}$ from $[m]$ uniformly *without* replacement to fill each column $k \in [s]$ of A, the fraction $\widetilde{q}_k$ of empty rows remaining after k rounds will be even smaller. Specifically, it holds

$$\mathbb{P}\left(\tilde{q}_k \leq t\right) \geq \mathbb{P}\left(q_k \leq t\right), \qquad \forall t \geq 0.$$

In turn,

$$\mathbb{P}\left(\forall k \in [s] : \tilde{q}_k \leq \exp(-k/s) + \varepsilon\sqrt{k/s}\right) \geq 1 - 2s\exp(-\varepsilon^2 m/2). \tag{16}$$

We now assume that the realization of the random matrix A yields, for each $k \in [s]$, $\tilde{q}_k \leq \exp(-k/s) + \varepsilon\sqrt{k/s}$. By the above, this occurs with probability exceeding $1 - 2s\exp(-\varepsilon^2 m/2)$. Continuing, let f_k be the fraction over m among the first dk balls thrown which form a *collision*, where a ball forms a collision if it lands in a bin which contains a ball thrown from a previous round. Note that $f_k = \tilde{q}_k - (1 - \frac{k}{s})$, the difference between the true fraction of empty bins and the fraction of bins that would be empty if there were no collisions. It follows that

$$\begin{aligned} f_k \leq \tilde{q}_k - 1 + \frac{k}{s} &\leq \exp(-k/s) - 1 + k/s + \varepsilon\sqrt{k/s} \\ &\leq \frac{k}{s}e^{-1} + \varepsilon\sqrt{k/s}, \end{aligned} \tag{17}$$

the last inequality holding because $\exp(-u) - 1 + u \leq \exp(-1)u$ for $u \in [0, 1]$.

Recall that $A \in \{0,1\}^{m \times s}$ is such that $a_{j,k} = 1$ if a ball is thrown into bin j during the kth round, and $a_{j,k} = 0$ otherwise. For $k \in [s]$, let $E_k \subset [m] \times [k]$ denote the subset of $|E_k| \leq dk$ entries such that $a_{j,\ell} = 1$, and let $C_k \subset E_k$ denote the subset of those entries corresponding to collisions. Then $|C_k| \leq dkf_k \leq dk(e^{-1} + \varepsilon)$. Recall now the assumption that $\boldsymbol{z} \in \mathbb{R}^s$ supported on $S = \{1, 2, \ldots, s\}$ is in decreasing rearrangement: $|z_1| \geq |z_2| \geq \cdots \geq |z_s|$. Observe that we can write

$$\sum_{(j,k) \in C_s} |z_k| \leq \sum_{k=1}^{s} n_k |z_k| \tag{18}$$

where $n_k \geq 0$ satisfies $\sum_{\ell=1}^{k} n_\ell = |C_k|$. Given the constraints $|z_1| \geq |z_2| \geq \cdots \geq |z_s|$ and $|C_k| \leq dk(e^{-1} + \varepsilon)$, the RHS expression of (18) is maximized by setting $n_k = d(e^{-1} + \varepsilon)$. Hence,

$$\sum_{(j,k) \in C_s} |z_k| \leq d(e^{-1} + \varepsilon)\|\boldsymbol{z}\|_1. \tag{19}$$

Since A has d ones per column by construction,

$$\sum_{(j,k) \in C_s} |z_k| + \sum_{(j,k) \in E_s \setminus C_s} |z_k| = \sum_{(j,k) \in E_s} |z_k| = d\|\boldsymbol{z}\|_1.$$

Combined with (19),

$$\sum_{(j,k) \in E_s \setminus C_s} |z_k| \geq d(1 - (e^{-1} + \varepsilon))\|\boldsymbol{z}\|_1,$$

and so

$$\|A\boldsymbol{z}\|_1 \geq \sum_{(j,k) \in E_s \setminus C_s} |z_k| - \sum_{(j,k) \in C_s} |z_k| \geq d(1 - 2(e^{-1} + \varepsilon))\|\boldsymbol{z}\|_1$$

and

$$\sum_{j=1}^{m}\left(\sum_{k=1}^{s} a_{j,k}|z_k|^2\right)^{1/2} \geq \sum_{(j,k)\in E_s\setminus C_s} |z_k| \geq d(1-(e^{-1}+\varepsilon))\|\boldsymbol{z}\|_1. \tag{20}$$

This finishes the proof. □

Applying a union bound over the $\binom{n}{s} \leq (n/s)^s$ subsets $S \subset [n]$ of size $|S| = s$, and over the $s! \leq s^s$ orderings of indices within any particular such subset S, Proposition 12 gives rise to a uniform result holding over all s-sparse vectors:

Corollary 13 (Corollary to Proposition 12)**.** *Fix $n, s, \in \mathbb{N}$, $\xi \in (0,1)$, and $\varepsilon > 0$. Fix $d \in \mathbb{N}$ satisfying*

$$d \geq 2\varepsilon^{-2}\log(n/\xi)$$

and let $m = ds \geq 2s\varepsilon^{-2}\log(n/\xi)$. Draw an $m \times n$ binary random matrix $A = (a_{j,k})$ with d ones per column chosen uniformly without replacement. With probability exceeding $1-\xi$ it holds that

$$d(1-2e^{-1}-2\varepsilon)\|\boldsymbol{z}\|_1 \leq \|A\boldsymbol{z}\|_1 \leq d\|\boldsymbol{z}\|_1 \qquad \forall \boldsymbol{z} \in \mathbb{R}^n : \|\boldsymbol{z}\|_0 \leq s \tag{21}$$

and

$$d(1-e^{-1}-\varepsilon)\|\boldsymbol{z}\|_1 \leq \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k}z_k^2\right)^{1/2} \qquad \forall \boldsymbol{z} \in \mathbb{R}^n : \|\boldsymbol{z}\|_0 \leq s$$

The inequalities in (21) imply that A/d satisfies the 1-restricted isometry property (1) of order s and level $\theta = 2e^{-1}+2\varepsilon$. Our probabilistic construction differs from existing constructions of 1-RIP matrices [5, 12] in that we use $m = sd$ rows but do not seek θ arbitrarily small, as opposed to using $m \geq \theta^{-1}sd$ rows to achieve 1-RIP for arbitrarily small $\theta > 0$.
Finally, we show that the results of Proposition 12 holding on each of the blocks $\boldsymbol{x}_{S_\ell}$ in the block-decreasing rearrangement of a vector $\boldsymbol{x}$ implies a lower bound on $\mathbb{E}_G\|(A_s \circ G)\boldsymbol{x}\|_1$ in terms of $\|\boldsymbol{x}\|_{1,2,s}$.

Proposition 14. *Consider $\boldsymbol{x} \in \mathbb{R}^n$ with s-block decreasing rearrangement $\boldsymbol{x} = (\boldsymbol{x}_{S_1}, \boldsymbol{x}_{S_2}, \ldots, \boldsymbol{x}_{S_{\lceil n/s\rceil}})$. Suppose, for some parameter $\gamma \in [0,1]$, that $A \in \mathbb{R}^{m\times n}$ satisfies*

$$d(1-\gamma)\|\boldsymbol{x}_{S_\ell}\|_1 \leq \sum_{j=1}^{m}\left(\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2}, \quad \ell = 1,2,\ldots,\lceil n/s\rceil. \tag{22}$$

Then

$$d(1-\gamma)\|\boldsymbol{x}\|_{1,2,s} \leq \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k}x_k^2\right)^{1/2}.$$

Proof. Consider the vector $\boldsymbol{y}_j = (y_{j,\ell}) \in \mathbb{R}^{\lceil n/s\rceil}$ with coordinates $y_{j,\ell} = \left(\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2}$. By the triangle inequality,

$$\|\sum_{j=1}^{m}\boldsymbol{y}_j\|_2 = \sqrt{\sum_{\ell=1}^{\lceil n/s\rceil}\left(\sum_{j=1}^{m}\left(\sum_{k\in S_\ell} a_{j,k}x_k^2\right)^{1/2}\right)^2} \leq \sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k}x_k^2\right)^{1/2} = \sum_{j=1}^{m}\|\boldsymbol{y}_j\|_2$$

Incorporating the assumptions (22) and recalling that $\|\boldsymbol{x}\|_{1,2,s} = \left(\sum_{\ell=1}^{\lceil n/s\rceil}\|\boldsymbol{x}_{S_\ell}\|_1^2\right)^{1/2}$ gives the desired result. □

3 Proof of Theorem 2

In this section we put together the ingredients to prove Theorem 2. Fix $\boldsymbol{x} \in \mathbb{R}^n$ with s-block decreasing rearrangement $\boldsymbol{x} = (\boldsymbol{x}_{S_1}, \boldsymbol{x}_{S_2}, \ldots, \boldsymbol{x}_{S_{\lceil n/s \rceil}})$. Fix parameter $\varepsilon > 0$. Consider $A = (a_{j,k}) \in \{0,1\}^{m \times n}$ a random binary matrix having d ones per column as constructed in Proposition 12. From that proposition,

$$\begin{aligned}
&\mathbb{P}\left(\exists \ell \in \{1,2,\ldots,\lceil n/s \rceil\} : \;\; d(1-e^{-1}-\varepsilon/2)\|\boldsymbol{x}_{S_\ell}\|_1 > \sum_{j=1}^m \left(\sum_{k \in S_\ell} a_{j,k} x_k^2\right)^{1/2}\right) \\
&\leq \sum_{\ell=1}^{\lceil n/s \rceil} \mathbb{P}\left(d(1-e^{-1}-\varepsilon/2)\|\boldsymbol{x}_{S_\ell}\|_1 > \sum_{j=1}^m \left(\sum_{k \in S_\ell} a_{j,k} x_k^2\right)^{1/2}\right) \\
&\leq 2(n/s+1)s\exp(-\varepsilon^2 m/8) \\
&\leq 4n\exp(-\varepsilon^2 m/8). \qquad (23)
\end{aligned}$$

We now assume that the realization of the random matrix A yields

$$d(1-e^{-1}-\varepsilon/2)\|\boldsymbol{x}_{S_\ell}\|_1 \leq \sum_{j=1}^m \left(\sum_{k \in S_\ell} a_{j,k} x_k^2\right)^{1/2}, \quad \ell = 1,2,\ldots,\lceil n/s \rceil,$$

and hence, by Proposition 14,

$$d(1-e^{-1}-\varepsilon/2)\|\boldsymbol{x}\|_{1,2,s} \leq \sum_{j=1}^m \left(\sum_{k=1}^n a_{j,k} x_k^2\right)^{1/2}.$$

By the above, this occurs with probability exceeding $1 - 4n\exp(-\varepsilon^2 m/8)$.

Consider now $G \in \mathbb{R}^{m \times n}$ having i.i.d. standard Gaussian entries, and let $\Phi = A \circ G$ where $\circ$ denotes the Hadamard (entrywise) product. By Lemma 6, $\mathbb{E}_G\|\Phi\boldsymbol{x}\|_1 = \beta_0 \sum_{j=1}^m \left(\sum_{k=1}^n a_{j,k} x_k^2\right)^{1/2}$ where $\beta_0 = \sqrt{2/\pi}$, and so by the above analysis it follows

$$\beta_0 d(1-e^{-1}-\varepsilon/2)\|\boldsymbol{x}\|_{1,2,s} \leq \mathbb{E}_G\|\Phi\boldsymbol{x}\|_1.$$

By Proposition 9, we have also the upper bound

$$\mathbb{E}_G\|\Phi\boldsymbol{x}\|_1 \leq 1.63\beta_0 d\|\boldsymbol{x}\|_{1,2,s}. \qquad (24)$$

Now, Proposition 8 gives that with respect to the draw of G,

$$\mathbb{P}\left(|\|\Phi\boldsymbol{x}\|_1 - \mathbb{E}_G\|\Phi\boldsymbol{x}\|_1| \geq \frac{\varepsilon}{2}\beta_0 d\|\boldsymbol{x}\|_{1,2,s}\right) \leq 2\exp\left(\frac{-\varepsilon^2\beta_0^2\|\boldsymbol{x}\|_{1,2,s}^2 d}{8\|\boldsymbol{x}\|_2^2}\right). \qquad (25)$$

We now assume that the realization of the random matrix G yields

$$|\|\Phi\boldsymbol{x}\|_1 - \mathbb{E}_G\|\Phi\boldsymbol{x}\|_1| \leq \frac{\varepsilon}{2}\beta_0 d\|\boldsymbol{x}\|_{1,2,s},$$

which by the above occurs with probability exceeding $1-2\exp\left(\frac{-\varepsilon^2\beta_0^2\|\boldsymbol{x}\|_{1,2,s}^2 d}{8\|\boldsymbol{x}\|_2^2}\right)$. Adding together the probabilities that either our assumption on A or our assumption on G does not hold, we have shown that with probability exceeding $1 - 4n\exp(-\varepsilon^2 m/8) - 2\exp\left(\frac{-\varepsilon^2\beta_0^2 m}{8s}\left(\frac{\|\boldsymbol{x}\|_{1,2,s}}{\|\boldsymbol{x}\|_2}\right)^2\right)$,

$$d\beta_0(1-e^{-1}-\varepsilon)\|\boldsymbol{x}\|_{1,2,s} \leq \|\Phi\boldsymbol{x}\|_1 \leq d\beta_0(1.63+\varepsilon)\|\boldsymbol{x}\|_{1,2,s}.$$

Setting $\Psi_s = \frac{1}{d\beta_0}(A \circ G)$, recalling that $d = m/s$, and using the bound $.63 \leq 1 - e^{-1}$, we recover the content of Theorem 2.

4 Discussion

Theorem 2 introduces a family of maps, which are shown to map finite-dimensional spaces equipped with the block ℓ_1/ℓ_2-norm to lower dimensional spaces equipped with the the ℓ_1-norm, while with high probability preserving the norm up to a constant distortion factor. In Euclidean space, by contrast, Johnson-Lindenstrauss embeddings can be made to have distortion arbitrarily close to 1, that is, there is only a factor that can be made arbitrarily small by increasing the embedding dimension.

In this section we show that some distortion factor is indeed necessary; we give two explicit families of examples in arbitrarily large dimensions for which the fraction of their block ℓ_1/ℓ_2-norm and the ℓ_1-norm of their image behave differently even in the asymptotic limit. Conditioned on the (dependent) random variables $a_{j,k}$, the entries of Φ_s and hence also the entries of $\Phi_s \boldsymbol{x}$ are independent Gaussian random variables. Thus the concentration of $\|\Phi_s \boldsymbol{x}\|_1$ around its mean is precisely understood, and it remains to compare the behavior of the mean for different instances of $\boldsymbol{x}$. As shown in Lemma 6 above, this boils down to studying the quantity $\sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2\right)^{1/2}$. The following proposition indeed provides two vectors with significantly different behavior of this quantity as normalized by the block ℓ_1/ℓ_2-norm.

Proposition 15. *Choose $a_{j,k}$ as defined in Theorem 2 and assume that $m \leq n$. Then for each $\eta, \nu > 0$, there exists a constant $C > 0$ such that if $s > C$ and $\frac{n}{s \log(n)} > C$, the following holds with probability at least $1 - \nu$.*

Consider $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n+1}$ with

$$\boldsymbol{x} = \left(1, \frac{1}{\sqrt{ns}}, \frac{1}{\sqrt{ns}}, \ldots, \frac{1}{\sqrt{ns}}\right)^t \text{ and } \boldsymbol{y} = (1, 0, \ldots, 0)^t.$$

Then one has $\frac{\sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k} x_k^2\right)^{1/2}}{d\|\boldsymbol{x}\|_{1,2,s}} \geq \sqrt{2} - \eta$, while $\frac{\sum_{j=1}^{m}\left(\sum_{k=1}^{n} a_{j,k} y_k^2\right)^{1/2}}{d\|\boldsymbol{y}\|_{1,2,s}} \equiv 1$.

Proof. To facilitate the calculations, we assume that s divides $n+1$, if this is not the case, one obtains an only slightly changed result. We first determine the block ℓ_1/ℓ_2-norm of $\boldsymbol{x}$ and $\boldsymbol{y}$, obtaining $\|\boldsymbol{y}\|_{1,2,s}^2 = 1$ and

$$\|\boldsymbol{x}\|_{1,2,s}^2 = (1 + \frac{s-1}{\sqrt{sn}})^2 + (\frac{n+1}{s} - 1)\frac{s}{n} = 2 + 2\frac{s-1}{\sqrt{sn}} - \frac{1}{n} + \frac{1}{ns} \leq 2 + 2\sqrt{\frac{s}{n}}. \tag{26}$$

To estimate the numerators, we note that for j fixed, the $a_{j,k}$ are independent Bernoulli random variables with parameter $p = \frac{1}{s}$. Thus $\mathbb{E}\sum_{k=2}^{n+1} a_{j,k} x_k^2 = \frac{1}{s^2}$ and, for $\delta > 0$,

$$\mathbb{P}\left(\sum_{k=2}^{n+1} a_{j,k} x_k^2 < (1-\delta)\frac{1}{s^2}\right) = \mathbb{P}\left(\sum_{k=2}^{n+1} a_{j,k} < (1-\delta)\frac{n}{s}\right).$$

The latter is a large deviation probability for the binomial distribution. It can be bounded via the relative entropy between two biased coins

$$H(a,p) = a\log(\frac{a}{p}) + (1-a)\log(\frac{1-a}{1-p}). \tag{27}$$

Then (for example by Theorem 1 in [1], applied to the expression with the roles of ones and zeros exchanged), the large deviation probability is bounded by

$$\mathbb{P}\left(\sum_{k=2}^{n+1} a_{j,k} < (1-\delta)\frac{n}{s}\right) \leq \exp\left(-nH\left(\frac{1+\delta}{s}, \frac{1}{s}\right)\right) \leq \left(\frac{1}{1+\delta}\right)^{(1+\delta)n/s} = e^{-cn/s}, \tag{28}$$

where $c = (1+\delta)\log(1+\delta) > 0$. Now we know that $a_{j,1} = 1$ for exactly d randomly chosen values of j. For these values of j, one has

$$\mathbb{P}\left(|\sum_{k=1}^{n+1} a_{j,k} x_k^2 < 1 + (1-\delta)\tfrac{1}{s^2} \right) \leq e^{-cn/s},$$

while for the other values of j, one has

$$\mathbb{P}\left(\sum_{k=1}^{n+1} a_{j,k} x_k^2 < (1-\delta)\tfrac{1}{s^2} \right) \leq e^{-cn/s}.$$

Thus with probability at least $1 - 2me^{-cn/s}$, one has

$$\begin{aligned}\sum_{j=1}^{m} \left(\sum_{k=1}^{n+1} a_{j,k} x_k^2 \right)^{1/2} &\geq d\sqrt{1 + (1-\delta)\tfrac{1}{s^2}} + (m-d)\sqrt{1-\delta}\tfrac{1}{s} \\ &\geq d(1 + \sqrt{1-\delta} - \tfrac{1}{s}(\sqrt{1-\delta})) \\ &\geq d(2 - \delta - \tfrac{1}{s})\end{aligned}$$

and consequently

$$\frac{\sum_{j=1}^{m} \left(\sum_{k=1}^{n} a_{j,k} x_k^2 \right)^{1/2}}{d\|\boldsymbol{x}\|_{1,2,s}} \geq \frac{2 - \delta - \frac{1}{s}}{\sqrt{2 + 2\sqrt{s/n}}}.$$

It is clear that if s and $\frac{n}{s}$ are large enough and if δ is chosen small enough, then this expression is ensured to get arbitrarily close to $\sqrt{2}$, as desired. On the other hand, as $m \leq n$, the associated probability of failure is bounded by $me^{-cn/s} \leq e^{\log(n) - cn/s}$, which, for fixed c and s, becomes arbitrarily small for n large enough.

The estimate for $\boldsymbol{y}$ follows directly from the fact that $a_{j,1} = 1$ for exactly d values of j and 0 otherwise. □

Remark 16. *The definition of $\boldsymbol{x}$ is inspired by the proof of Proposition 9. Namely, the first inequality in* (8) *is sharp when the two terms are equal, and the last inequality in* (10) *is the sharper, the more the two summands differ. So the constant in Proposition 9 can be significantly improved unless both of these facts happen at the same time. This basically boils down to having a jump within the first block and roughly the same block norm contribution of the first block and the tail, which $\boldsymbol{x}$ is an extreme example of.*

Proposition 15 provides a counterexample, for which a constant distortion is necessary in Theorem 2. Considering that our proofs above yield the exact analogue to Theorem 2 also for the interpolation norm $K(\boldsymbol{x}, \sqrt{s})$ instead of the block norm $\|\boldsymbol{x}\|_{1,2,s}$, one may ask whether the former norm has better empirical performance. To answer this question, we numerically test the validity of Theorem 2 for these two norms on four different types of signals for $n = 1000$, $s = 10$ fixed and varying number of measurements. The results are presented in Figure 1; for each of the norms, we plot the relative distortion $\frac{\|\Psi_s \boldsymbol{x}\|_1 - \|\boldsymbol{x}\|}{\|\boldsymbol{x}\|}$, where Ψ_s is as in Theorem 2 and $\|\cdot\|$ is the norm in question.

In example (a), s-sparse signals are generated by choosing a support at random and the corresponding entries according to the standard normal distribution. In example (b), we consider the fixed 1-sparse vector supported in the first position. In example (c), we consider signals inspired by the counterexample of Proposition 15 with one large entry and the other entries chosen according to a normal distribution with considerably smaller variance. Finally,

in example (d), we consider signal with independent entries drawn from the standard normal distribution. Our experiments show that for both norms, the average ℓ_1-norm of the image of normalized vectors in the different classes behave differently. As expected, for the block norm, the signals inspired by the counterexample yield images with a larger ℓ_1-norm than for 1-sparse and also for random signals. On the other hand, for the interpolation norm, images of random signals typically have smaller norm, while the vectors inspired by $\boldsymbol{x}$ and $\boldsymbol{y}$ in Proposition 15 have a comparable behavior. These observations made us choose to present our results in terms of the block norm rather than the interpolation norm, as for the block norm, the upper and lower distortion factors essentially disappear for random signals (which we see as representing the generic behavior). Notably, in the example of s-sparse signals, the resulting images have a smaller norm. In this case, the behavior for the two norms is identical, as for sparse vectors they both reduce to the ℓ_1-norm.

Acknowledgment

We thank Anna Gilbert, Arie Israel, Joe Neeman, Jelani Nelson, and Mark Rudelson for helpful input and suggestions. We also thank the anonymous referees for their valuable comments. F. Krahmer was supported by the German Science Foundation (DFG) in the context of the Emmy Noether Junior Research Group "RaSenQuaSI" (grant KR 4512/1-1). R. Ward was supported in part by an AFOSR Young Investigator Award, DOD-Navy grant N00014-12-1-0743, and an NSF CAREER Award.

References

[1] R. Arratia and L. Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of Mathematical Biology*, 51(1):125–131, 1989.

[2] R. Baraniuk and M. Wakin. Random projections of smooth manifolds. In *Found. Comp. Math*, pages 941–944, 2006.

[3] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. Wakin. A simple proof of the Restricted Isometry Property for random matrices. *Constr. Approx.*, 28(3):253–263, 2008.

[4] C. Bennett and R. Sharpley. *Interpolation of operators*. New York, 1988.

[5] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 798–805, 2008.

[6] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft. When is "nearest neighbor" meaningful? *Proc. Internat. Conf. Database Theory*, pages 217–235, 1999.

[7] T. T. Cai, L. Wang, and G. Xu. New bounds for restricted isometry constants. *IEEE Trans. Inform. Theory*, 56(9):4388–4394, 2010.

[8] E. Candès, Y. Eldar, D. Needell, and R. Paige. Compressed sensing with coherent and redundant dictionaries. *Appl. Comput. Harmon. Anal.*, 31:59–73, 2011.

[9] M. Charikar and A. Sahai. Dimension reduction in the ℓ_1 norm. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 551–560. IEEE, 2002.

[10] Z. Cvetkovski. Schur's inequality, Muirhead's inequality and Karamata's inequality. In *Inequalities*, pages 121–132. Springer Berlin Heidelberg, 2012.

[11] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

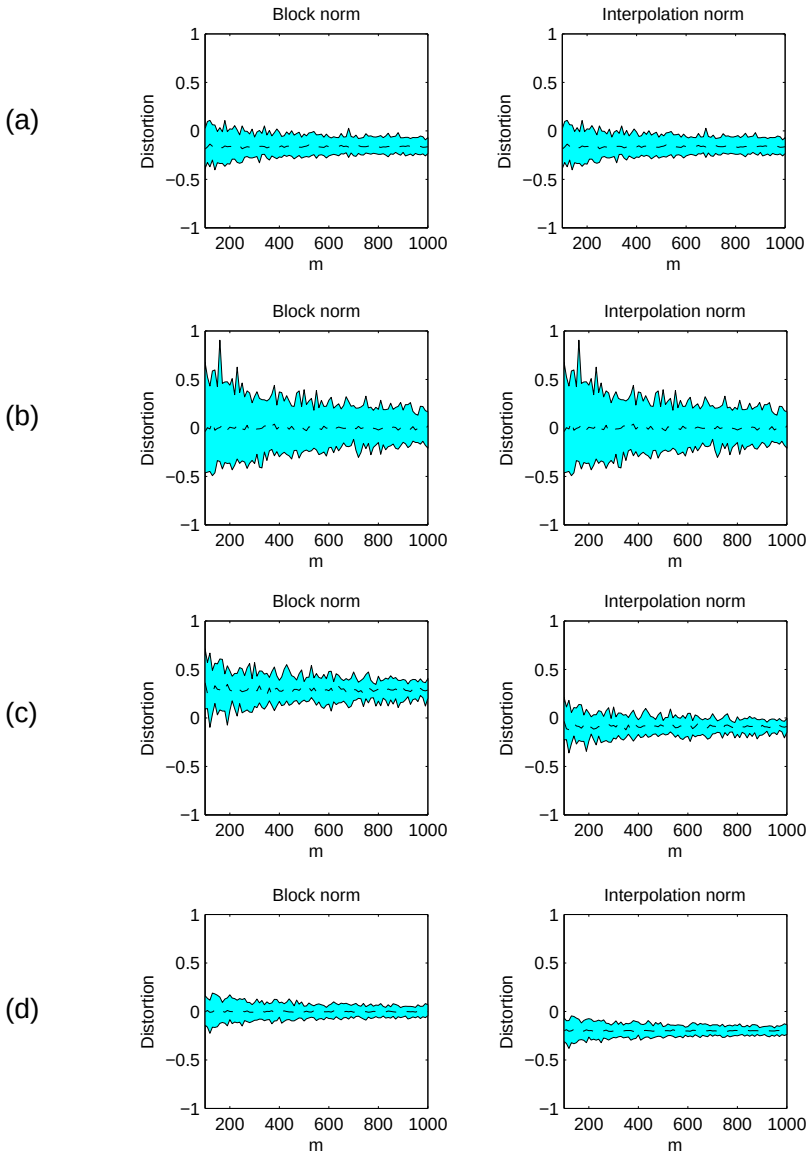

Figure 1: A comparison of the distortion $(\|\Psi_s(x)\|_1 - \|x\|)/\|x\|$ where $\|x\|$ is either the block norm $\|x\|_{1,2,s}$ or the interpolation norm $\mathcal{K}(x, \sqrt{s})$ with Ψ_s as in Theorem 2 and $s = 10$ fixed throughout. We consider distortion statistics (the minimum, maximum, and median) over 64 instances of Ψ_s over several signal classes: (a) s-sparse signals, (b) 1-sparse signals, (c) a two-level signal inspired by Proposition 15, and (d) Gaussian random signals.

[12] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.

[13] A. Y. Garnaev and E. D. Gluskin. The widths of a Euclidean ball (Russian). *Dokl. Akad.*

Nauk SSSR, 277(5):1048–1052, 1984.

[14] U. Haagerup. The best constants in the khintchine inequality. *Studia Math.*, 70(3):231–283, 1981.

[15] N. Halko, P. Martinsson, and J. Tropp. Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53.2:217–288, 2011.

[16] A. Hinneburg, C. Aggarwal, and D. Keim. What is the nearest neighbor in high dimensional spaces? *in Proc. 26th Internat. Conf. Very Large Data Bases*, pages 506–515, 2000.

[17] T. Holmstedt. Interpolation of quasi-normed spaces. *Math. Scand.*, 26:177–199, 1970.

[18] P. Indyk. Algorithmic applications of low-distortion embeddings. *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[19] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. *J. ACM*, 53(3):307–323, 2006.

[20] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math*, 26:189–206, 1984.

[21] D. Kane, J. Nelson, and D. Woodruff. On the exact space complexity of sketching and streaming small norms. *SODA*, pages 1161–1178, 2010.

[22] F. Krahmer and R. Ward. New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property. *SIAM Journal on Mathematical Analysis*, 43(3):1269–1281, 2011.

[23] J. Lee, M. Mendel, and A. Naor. Metric structures in l_1: Dimension, snowflakes, and average distortion. *European Journal of Combinatorics*, 26(8):1180–1190, 2005.

[24] P. Li. Estimators and tail bounds for dimension reduction in ℓ_q, $(0 < q \leq 2)$ using stable random projections. *SODA*, pages 10–19, 2008.

[25] E. Liberty, F. Woolfe, P. Martinsson, V. Rokhlin, and M. Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.

[26] A. E. Litvak, A. Pajor, M. Rudelson, and N. Tomczak-Jaegermann. Smallest singular value of random matrices and geometry of random polytopes. *Advances in Mathematics*, 195(2):491–523, 2005.

[27] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.

[28] S. J. Montgomery-Smith. The distribution of Rademacher sums. *Proceedings of the American Mathematical Society*, 109(2):517–522, 1990.

[29] J. Nelson and D. Woodruff. Fast Manhattan sketches in data streams. *PODS*, pages 99–110, 2010.

[30] I. Newman and Y. Rabinovich. Finite volume spaces and sparsification. *arXiv preprint arXiv:1002.3541*, 2010.

[31] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 2013.

[32] T. Sarlos. Improved approximation algorithms for large matrices via random projections. *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[33] G. Schechtman. Dimension reduction in l_p, $0 < p < 2$. *arXiv preprint arXiv:1110.2148*, 2011.

[34] M. Thorup and Y. Zhang. Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation. *SIAM J. Comput.*, 41(2):293–331, 2012.

[35] R. Ward. Compressed sensing with cross validation. *IEEE Trans. Inform. Theory*, 55:5773–5782, 2009.

5 Appendix

5.1 Verifying that $\|\cdot\|_{1,2,s}$ is a norm

Here we verify that the function $\|\boldsymbol{x}\|_{1,2,s} := \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{S_\ell}\|_1^2}$ is indeed a norm.

It is straightforward that $\|\boldsymbol{x}\|_{1,2,s} \geq 0$ and that $\|\boldsymbol{x}\|_{1,2,s} = 0$ implies $\boldsymbol{x} = 0$. It is also clear that $\|a\boldsymbol{x}\|_{1,2,s} = |a|\|\boldsymbol{x}\|_{1,2,s}$. It remains to verify the triangle inequality: $\|\boldsymbol{x}+\boldsymbol{y}\|_{1,2,s} \leq \|\boldsymbol{x}\|_{1,2,s} + \|\boldsymbol{y}\|_{1,2,s}$ for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$. To do this, let us set up some notation. Denote the partitioned supports in the block decreasing rearrangement of $\boldsymbol{x}$ by $S_1, S_2, \ldots$, the partitioned supports in the block decreasing rearrangement of $\boldsymbol{y}$ by $T_1, T_2, \ldots$ and the partitioned supports of the block decreasing rearrangement of $\boldsymbol{x}+\boldsymbol{y}$ by $U_1, U_2, \ldots$. Then

$$\|\boldsymbol{x}+\boldsymbol{y}\|_{1,2,s} = \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|(\boldsymbol{x}+\boldsymbol{y})_{U_\ell}\|_1^2} \leq \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{U_\ell}\|_1^2} + \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{y}_{U_\ell}\|_1^2}$$

thanks to the block ℓ_1/ℓ_2 vector norm (with support sets fixed) satisfying the triangle inequality. Now, $\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{U_\ell}\|_1^2 \leq \sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{S_\ell}\|_1^2$ can be seen to hold by appealing to *Karamata's inequality* [10] to the sequences $\boldsymbol{r}_1 = (\|\boldsymbol{x}_{S_1}\|_1, \|\boldsymbol{x}_{S_2}\|_1, \ldots)$ and $\boldsymbol{r}_2 = (\|\boldsymbol{x}_{U_1}\|_1, \|\boldsymbol{x}_{U_2}\|_1, \ldots)$, noting that $\boldsymbol{r}_1$ majorizes $\boldsymbol{r}_2$. Using the same argument to show $\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{y}_{U_\ell}\|_1^2 \leq \sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{y}_{T_\ell}\|_1^2$, we then have that the RHS expression above is

$$\leq \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{x}_{S_\ell}\|_1^2} + \sqrt{\sum_{\ell=1}^{\lceil n/s \rceil} \|\boldsymbol{y}_{T_\ell}\|_1^2} = \|\boldsymbol{x}\|_{1,2,s} + \|\boldsymbol{y}\|_{1,2,s}, \tag{29}$$

verifying the triangle inequality.

5.2 Proof of Lemma 7

Recall that $Y_i \sim \mathcal{N}(0, \sigma_i^2)$ has density function is $dF_i(t) = \beta_0 \frac{1}{2\sigma_i} \exp\left(-\frac{t^2}{2\sigma_i^2}\right)dt$ where $\beta_0 = \sqrt{2/\pi}$. We may then estimate for each $u \geq 0$

$$\begin{aligned}
\mathbb{E}\left[\exp\left(u(|Y_i| - \beta_0\sigma_i)\right] &= \int_{t=-\infty}^{\infty} \exp\left(u(|t| - \beta_0\sigma_i)\right) dF_i(t) \\
&= \frac{\beta_0}{\sigma_i} \int_{t=0}^{\infty} \exp(u(t - \beta_0\sigma_i)) \exp\left(-\frac{t^2}{2\sigma_i^2}\right) dt \\
\text{(set } s = \sqrt{2}t/\sigma_i - \sqrt{2}\sigma_i u\text{)} \qquad &= \exp\left(u^2\sigma_i^2/2 - u\beta_0\sigma_i\right)\sqrt{2}\left[(\beta_0/2)\int_{s=-\sqrt{2}\sigma_i u}^{\infty} \exp\left(-s^2/4\right)ds\right] \\
&= \exp\left(u^2\sigma_i^2/2 - u\beta_0\sigma_i\right)\sqrt{2}\left[1/2 + (\beta_0/2)\int_{s=0}^{\sqrt{2}\sigma_i u} \exp\left(-s^2/4\right)ds\right]
\end{aligned}$$

One of the two cases holds:

1. If $u \geq \frac{\log(2)}{2\beta_0\sigma_i}$, then $\sqrt{2} \leq \exp(u\beta_0\sigma_i)$

2. If $0 \leq u < \frac{\log(2)}{2\beta_0\sigma_i}$, then

$$\sqrt{2}\left[1/2 + (\beta_0/2)\int_{s=0}^{\sqrt{2}\sigma_i u} \exp\left(-s^2/4\right)ds\right] \leq 1/\sqrt{2} + \frac{\beta_0}{\sqrt{2}}\sqrt{2}\sigma_i u \leq 1$$

In either case, we may bound the final RHS expression above to estimate

$$\mathbb{E}\left[\exp\left(u(|Y_i| - \beta_0\sigma_i)\right] \leq \exp\left(u^2\sigma_i^2/2\right). \tag{30}$$

A similar analysis reveals that also

$$\mathbb{E}\left[\exp\left(-u(|Y_i| - \beta_0\sigma_i)\right] \leq \exp\left(u^2\sigma_i^2/2\right). \tag{31}$$

Recall that the random variable of interest is of the form $Z = \frac{1}{d}\sum_i(|Y_i| - \mathbb{E}|Y_i|) = \frac{1}{d}\sum_i(|Y_i| - \beta_0\sigma_i)$ with $Y_i \sim \mathcal{N}(0, \sigma_i^2)$ independent. Recall also that $\sum_i \sigma_i^2 \leq d\alpha^2$ by assumption. It follows that

$$\begin{aligned}
\mathbb{E}(\exp(uZ)) &= \mathbb{E}\exp\left(\frac{u}{d}\sum_i(|Y_i| - \beta_0\sigma_i)\right) \\
&= \mathbb{E}\prod_{i=1}^m \exp\left(\frac{u}{d}(|Y_i| - \beta_0\sigma_i)\right) \\
&= \prod_{i=1}^m \mathbb{E}\left[\exp\left(\frac{u}{d}(|Y_i| - \beta_0\sigma_i)\right)\right] \\
&\leq \prod_{i=1}^m \exp\left(\frac{u^2\sigma_i^2}{2d^2}\right) \\
&= \exp\left(\frac{u^2}{2d^2}\sum_i \sigma_i^2\right) \leq \exp\left(\frac{u^2\alpha^2}{2d}\right)
\end{aligned}$$

Recall that if a random variable X satisfies $\mathbb{E}[X] = 0$ and $\mathbb{E}\exp(uX) \leq \exp(Cu^2)$ for all $u \in \mathbb{R}$ and for some constant $C > 0$, then $\mathbb{P}(|X| \geq \lambda) \leq 2\exp(-\frac{\lambda^2}{4C})$ for each $\lambda \geq 0$ (see, for example, Proposition 7.24 of [12]). Since $\mathbb{E}[Z] = 0$, it follows that $\mathbb{P}(|Z| \geq \lambda) \leq 2\exp(-\frac{\lambda^2 d}{2\alpha^2})$ for each $\lambda \geq 0$. This proves Lemma 7.